I have had my solutions for a long time,
but I do not yet know how I am to arrive at them.

-C. F. Gauss

Chapter 6

Conclusions and Future Directions

This dissertation has been a study of the statistics of scale space. It not only includes a statistical analysis of current methods in multiscale differential geometry for image analysis, but it also presents new multiscale statistical methods for use with multivalued data. The first sections of this chapter summarize the contributions and results of dissertation. Other sections of this chapter foreshadow future work in the theory of multiscale image statistics and applications for this research.

In an effort to automatically control nonlinear diffusion for image segmentation, I began this research with an analysis of the effects of noise within the image on multiscale partial derivative measurements. The result was a theoretical study relating noise, scale, and multiscale differential invariants. The contributions of this study to the field of image analysis are listed in the next section.

By themselves, however, the relations between geometry, scale, and noise derived through this analysis were not sufficient to automatically select control parameters for nonlinear diffusion systems, particularly those involving multiple intensities in each pixel. To address the problems of incommensurable image values in multiparameter data, I have introduced a new form of scale-space analysis for image analysis, suggestive of a new approach to multivalued image segmentation. This scale-space analysis approach is based on a form of statistical moments that provide multiscale statistical invariance. Like the family of differential invariants, these moments can be made invariant to spatial transformations such as rotation, translation, and zoom (the simultaneous scaling of image magnification and measurement aperture). These statistical measurements are also easily made invariant with respect to linear functions of intensity. This dissertation has explored some of the statistical properties of scale-space differential invariants, and it has described new statistical operators that are comparable in function to differential invariants. These contributions are also summarized in Section 6.1.

To bring this dissertation past providing a theoretical structure, work must be done to explore the power, the properties, and the limitations of multiscale image statistics. Applications proving the capabilities enabled by statistically based methods are essential next steps. Examples have been provided in the text to demonstrate the effectiveness of these methods and their potential. However, while some scale-space applications have been shown, the behavior of multiscale statistics across scale-space remains a substantially unexplored area. Much work remains not only to validate the use of statistical moments, but also to compare their use to other methods, in particular the set of scale-space differential operators. Finally, in support of this thesis, multiscale statistics should be used to solve real problems. Future directions for this research are outlined later in this chapter. Possible applications are suggested that will emphasize the role of multiscale statistics.

6.1. Contributions and Conclusions

This dissertation has included three substantive developments as original work. The first development was an analysis of error propagation from spatially uncorrelated noise in an image to multiscale differential invariant measures of image structure. The second original development has been multiscale image analysis based on local statistics rather than on local partial derivatives as geometric measurements. The third contribution has been the development of directional multiscale image statistics.

The first of these developments, presented in Chapter 3, involved a new study on the effects of spatially uncorrelated noise on normalized linear scale space measurements. The presentation examined the current model of normalized scale-space differential invariants and generated mathematical expressions for the propagation of noise. To allow differential invariants to be used in scale space, ordinary derivatives at scale must be normalized by the size of their measurement aperture. The resulting dimensionless quantities allow comparisons between different multiscale partial derivative operators across varying orders of differentiation. Surprisingly, results showed that for all given levels of initial intensity of noise, the absolute error in the multiscale derivative decreases between zeroth and first order measurements. Moreover, though the level of propagated noise increases thereafter with increasing order of differentiation, it remains less than the initial error until the third or fourth order derivatives are taken. This finding brings into some question the common wisdom that low order differentiation badly propagates noise.

New image analysis methods based on multiscale statistical invariants and their implicit scale spaces were developed in Chapter 4. Local, isotropic, Gaussian neighborhood sampling operators of varying scale were used to generate local central moments of intensity that capture information about the local probability distribution of intensities at a pixel location under the assumption of piecewise ergodicity. Image analysis based on multiscale image statistics can easily be made invariant to linear functions of intensity as well as spatial rotation and translation. This trait makes the identification of objects independent of the absolute brightness of the object as well as independent of its orientation and position within the image. Multiscale central moments of intensity of different order demonstrated properties similar to multiscale differential geometric operators. Specifically, multiscale variances reflect boundariness in a fashion similar to the multiscale gradient magnitude, and multiscale skewness shows a response similar to the multiscale Laplacian operator. A scale-space algorithm for selecting locally adapted normalizing variance measures for variable conductance diffusion was developed based on these moments. Also, multiscale statistics have been generalized to provide a means of analyzing multivalued data where the data channels within the image are incommensurable (i.e., they have no common metric for measurement).

The spatially isotropic operators from Chapter 4 did not adequately capture image structure. The underlying geometry of the image introduces biases in the measurement of the distribution of noise within the system using isotropic operators. Orientable operators that were introduced in Chapter 5 allowed central moments to be measured in preferred directions. These central moments reduced bias in noise characterization that arises from the gradient of the image function. I showed how singular value decomposition (SVD) of directional covariances produced principal axes indicating maximum and minimum spread of the directional probability distributions. The derivation suggested that these orthogonal directions and their corresponding eigenvalues can be used to normalize measurements made of local intensities. Multivalued image analysis has also been shown to be possible through directional covariances. Using the generalized form of SVD called canonical analysis, normalized multilocal coordinate systems based on covariances enable relating the behavior of the various intensities where the lack of common metrics make analysis via intensity vectors inappropriate.

Before this research, processing of multivalued images have been based on *ad hoc* geometric principles, unlike scalar image processing. If the relationships between the separate channels of a multivalued image are known *a priori*, a mapping of the space spanned by the separate image values can provide a reparameterization of the values allowing relations to be drawn among them. Since information about the intervalue relations is normally absent, the assumption is implicitly made that the multiple values have a common basis for comparison and can be treated as vectors. Methods based on multiscale differential geometric measurements are applicable to scalar-valued images, but are not appropriate for multivalued images such as multi-echo sequences from a medical magnetic resonance scan. This dissertation has leveled that discrepancy by providing tools that by measuring correlations among multiple values of image data provide common metrics for comparison among them. The methods presented here, based on multiscale image statistics, provide the foundation from which to start building new theories and approaches to understanding multivalued data and analyzing and processing multivalued images.

6.2. Future Directions in Multiscale Statistical Theory

While the introduction and development of new multiscale statistical invariants has been a substantial task, there remain many unanswered questions. Fortunately, there are also many related fields from which to borrow analytical tools as well as motivations and ideas. Linear scale space and its differential invariants have been much of the model for this work. Pursuing parallels between scale-space geometry and scale-space statistics is likely to yield important and interesting results.

The field of statistics also presents a wide body of knowledge upon which future explorations into multiscale moments of image intensity can be based. This work entails the introduction and presentation of multiscale central moments of images. The uses of these moments are within the purview of applied statistics and computer science. In particular, the analysis and relationships among these moments is a problem in probability. Skilled probabilists may be interested in exploring these spatially windowed central moments, their properties, and their limitations.

I address some of the salient points of these issues in the following paragraphs. This discussion is not intended to be a complete survey but rather an introduction to a few research areas with some insights into promising intellectual directions.

6.2.1. Local Differential Geometry and Local Image Statistics

There are many parallels between the study of scale-space differential invariants and multiscale image statistics. As the understanding of images through differential geometry advances, the understanding of scale spaces of multiscale image statistics should also grow. Questions that are raised about differential geometric operators are often also illuminating when directed toward multiscale statistics.

I have included some work on the propagation of noise in multiscale image statistics. The full characterization of the effects of noise in this area is needed. The results presented in Chapter 3 on linear scale-space differential operators indicate that there are some limitations on how many orders of differentiation are expected to be useful in normalized scale-space image analysis. It remains to be seen if parallel findings will be found of multiscale central moments of image intensity.

Perhaps a more important question is in the propagation of noise through nonlinear scale. Since the Gaussian represents a solution to the heat equation, and given the extensive justification of the Gaussian as a preferred linear operator in image analysis, diffusion is an ideal process for studying linear scale space. There has been considerable work on nonlinear scale spaces, using nonlinear diffusion equations to process and analyze images [e.g., Florack 1993, Gerig 1992, Whitaker 1993ab]. Empirical evidence indicates that the variance of white noise in the input signal is reduced through many of these nonlinear diffusion systems. An exhaustive analysis, reapplying some of the approaches of Chapter 3, may prove fruitful and may induce advances in computer vision.

6.2.2. Multiscale Distribution Analysis

Beyond applying the ideas of differential geometry to multiscale statistical invariants, there are motivations and research directions suggested by extending statistical methods to multiscale central moments. At the turn of the century, Pearson developed a taxonomy for classifying a wide range of univariate probability distributions [see Johnson 1970]. Pearson’s system of distributions is based on the qualitative shape of probability density function as governed by four input parameters. These parameters are strongly related to the mean, the variance, the skewness, and the kurtosis of the distribution. The fourth parameter seems to have small influence on the classification. It was considered impractical to include a fifth or sixth parameter.

It will be interesting to explore local distributions and their properties in these terms. Using the results of Chapter 4, local measures of the mean and central moments of up to the fourth order have been presented in this dissertation. Pearson’s classification system could easily be applied to these moments, and a taxonomy of local distributions constructed. Such a multiscale classification system of local intensity distributions may illuminate issues of boundariness and medialness in image analysis.

6.2.3. Comparing Two Distributions

Qualities such as texture, anisotropy, boundariness, and medialness can be used to distinguish image segments. These qualities can measured through an analysis of the local probability distributions of multiscale image statistics. The resulting information could allow the comparison of two pixels based on the probability distribution exhibited by their local neighborhoods. A useful statistical means of comparing two distributions is the Kolmogorov-Smirnov statistic.

There are many statistical means of comparing two distributions to evaluate their similarity. For data that are discretized or "binned," the chi-square test is a powerful evaluation tool. For unbinned data that are functions of a single independent variable (e.g., two valued functions of time or space), the cumulative distribution function (or an unbiased estimator of it) $S_N(x)$ of data can be used in meaningful comparisons. Cumulative distributions of two separate random variables can then be compared. This process is captured by the Kolmogorov-Smirnov (K-S) statistic,

$$\mathrm{D} = \max_{-\infty < x < \infty} \left| S_{N_1}(x) - S_{N_2}(x) \right| \tag{6-1}$$

A complete description of the K-S statistic as well as source code implementing its computation have been published by Press, *et al.* [Press 1992].

There are many properties of the K-S statistic that make it useful in the context of multiscale image statistics. Finding the significance of the statistic is a tractable computation. Its invariance under certain transformations of intensity is also beneficial. The following is a quote from Press, *et al.*

> What makes the K-S statistic useful is that *its* distribution in the case of the null hypothesis (data sets drawn from the same distribution) can be calculated, at least to useful approximation, thus giving the significance of any observed nonzero value of *D*. A central feature of the K-S test is that it is invariant under reparameterization of *x*; … For example, you will get the same significance using *x* as log *x*. [Press 1992].

The K-S test can be recast using a Gaussian spatial windowing function to attain a spatially weighted distribution. Questions comparing such qualities as texture, anisotropy, boundariness, and medialness may be resolved through a local K-S test of multiscale image statistics.

6.3. Applying Multiscale Image Statistics

Promising directions for this work reside in application areas. Much of this research has been based on a particular model of images; images are assumed to have piecewise ergodic intensity distributions across space. Multiscale image statistics will enable new segmentation methods for these images, particularly those images containing multiple values per pixel.

The results from this dissertation enable multivalued nonlinear diffusion on images with incommensurable data values. Multimodal medical images, multivalued information from cartographic sources, and even color images provide sources for multiparameter data with possibly incommensurable values. Nonlinear diffusion controlled by multiscale directional image statistics is a likely means for performing smoothing of the intensity values while preserving the geometry of the image represented by object boundaries. Some of the ideas generated by this research are presented in this section indicating possible research directions.

Other important issues in segmentation include the difficulties of representing uncertainty in the composition of a pixel value. Such situations arise when a pixel represents a composite or mixture of semantic elements. This often occurs at object boundaries where a pixel may reside only partially within an object. The resulting state is not appropriate for a binary segmentation. However, it is easily represented in the form of probabilities. Multiscale image statistics provide a background and foundation to assign probability values to pixels as well as directing segmentation algorithms to likely orientations to assign connectivity relations among object segments.

6.3.1. Statistical Control of Nonlinear Diffusion

I have shown how multiscale image statistics can be used to set control parameters for boundary preserving noise smoothing systems based on nonlinear, geometry driven, or variable conductance diffusion (VCD). Section 4.8 describes methods to incorporate these measures directly into the diffusion equations, generating a new form of this type of image analysis. Practical studies demonstrating these methods on real data, both scalar-valued and multivalued, are required to measure their effectiveness in solving real problems.

This particular approach also needs to be refined in light of current developments. As shown by the work of Eberly [Eberly 1994] and supported by the canonical analysis arguments of Chapter 5, the gradient values of the conductance functions should be normalized not only by the local variance but also by the scale at which they are measured. Chapter 5 also introduces the acquisition of second moment values along the minor axis. These results are orthogonal to the variance measured in the direction of greatest change and are likely to be a less biased measure of the local noise. Isophote curvature will still introduce some bias, but an investigation of these properties is warranted given these dissertation results.

Gerig and Whitaker both have generalized some forms of VCD to higher dimensions. Gerig has demonstrated vector-valued nonlinear diffusion on medical images with some success [Gerig 1992]. Whitaker has shown that invariants other than zeroth order intensities can be diffused; his resulting geometry limited diffusion has been able to generate ridge structures that describe the general form of objects within images [Whitaker 1993ac].

While the resulting user-supervised statistically constrained VCD filtering mechanism provides a principled means of measuring dissimilarity or gradients of possibly incommensurable within-pixel data parameters, there remains the problems associated with the variations of local intensity and the non-stationary nature of intensity and contrast in MR images. The development of the local multivalued statistics from Chapters 4 and 5 address these issues. The insights presented in this thesis make possible more robust explorations into statistically controlled VCD methods on multivalued data. The eventual goal of automated VCD has become more likely through the statistical approaches presented in this dissertation.

Beyond controlling segmentation systems, multiscale image statistics may also play a significant role in understanding theoretical aspects of scale spaces. The use of multiscale directional statistics may be instrumental in the study of nonlinear scale spaces. Measuring the decreasing values of the minimum local directional variance at a pixel location may be an effective means of tracking object behavior across nonlinear scale. Nonlinear scale spaces will require new distance metrics; understanding the rules for these metrics may come from statistical analyses.

6.3.2. Mixtures in Segmentation

A single pixel often cannot be represented by a single segment. Rather it is a composite of different object types. This often happens when an object boundary intersects the area sampled by a pixel. It also often occurs when the image data represents fine structure that approaches the dimensions of a pixel. Such data include roadways on geographic maps or fine blood vessels in medical images. In such cases it is inappropriate to assign such a pixel to a single segment designating it as all one type of value. However, by casting the problem in statistical language, the problem is made approachable. Probabilities may be assigned to a pixel, reflecting the likelihood of multiple segments within that pixel's area. These probabilities can be used as a measure of the fraction of different image segments within a single pixel.

There are two approaches to assigning probabilities to handle the composite pixel problem. One approach is to use some outside means of generating an ensemble of images over which probability distributions can be estimated and probabilities assigned. Other information about the physics of the data acquisition method can be used to assist the segmentation. For example, in his dissertation Laidlaw calculates probability distributions at each pixel of an medical magnetic resonance image from which he generates a segmentation [Laidlaw 1995].

The other approach toward statistical segmentation involves modeling local image structure as ergodic. This has been presented in this dissertation. The ergodic assumption allows spatial sampling, rather than ensemble sampling, to generate information about the probability distribution of intensities at a pixel location since boundary properties and issues of object scale can be gracefully handled. These developments in multiscale central moments enable probability distributions to be calculated and multiple probabilities assigned to individual pixels. How these probabilities and composite pixels are combined with contiguous segments is a topic for future research. Also, combining ensemble statistics within a pixel with multiscale image statistics that capture probability distributions over image space is likely to be a productive research direction.

6.4. Multiscale Image Statistics in Medicine

The origins of this research are founded deep within applications of image processing methods in medicine. It is my desire to see this research transferred to the clinic where principled analytic methods can be used to provide answers and aids to diagnosis and treatment. It is essential that computer aids used in health care be robust, repeatable, verifiable, and understandable. Credibility and reliability are required attributes of any computer program in a health care situation.

The medical field in particular is a rich source for multiparameter image data where the separate within-pixel values are not commensurable. There is an oncoming flood of multimodality images based on the registering and fusing of intensity values of multiple images of the same anatomy, acquired through independent means. Combinations of nuclear medicine, computed tomography, magnetic resonance imaging, and/or electro-encephalogram information are beginning to inundate the field of medical image processing. These data contain multiple values which are inherently incommensurable. The multivalued statistical methods described in this dissertation should provide a principled foundation for the processing of these new forms of multimodality images.

This work is directed toward applications of magnetic resonance imaging, improving image quality and attempting to find approachable means to provide segmentation and classification in clinical images. The images generated by MRI often exhibit structured variations of intensity of low frequency, allowing correction along smooth, gentle gradients. I chose to explore this method as a means of seeking a correction mechanism for these non-homogenous nonlinear responses of tissue in MR images.

MRI has particularly nice attributes when viewed through the metaphor of multiscale image statistics. It is not only possible, but it is often the case that multiple values of the image are acquired. Figure 6.1 shows a two echo axial MR 2D image of the head along with its scatterplot histogram. These values cannot be considered vector quantities, since they are not tied to a common measurement metric. However, they do represent registered image intensity values, dependent on spatial location. There are strong correlations of image intensity within the pixel values, depending on the tissue they portray. Multivalued statistics can be applied to these images, and meaningful results extracted.

Figure 6.1. A 2D dual-echo MR image of the head with its scatterplot histogram.

Another useful feature of MR imaging is that issues of small sample size can be addressed through careful acquisition technique. The typical MR acquisition method attempts to improve signal to noise in diagnostic images by taking several samples or images of any single slice plane and averaging them together. Noise and registration artifact are thus reduced by the averaging of registered images. This is a simple method of using the statistics of an ensemble of images to improve the measurement of the mean intensity value. This approach is easily generalized to include statistics other than the mean, including the infinite set of central moments. With a sufficient number of samples in an ensemble of images, the probability distribution of intensities at every pixel location could be completely characterized.

The practical aspects of VCD in medicine are being explored at many institutions, where non-linear diffusion filters often serve a pre-processing role before traditional statistical pattern recognition classifiers are applied. In particular, filtering mechanisms provided by Guido Gerig [Gerig 1991, 1992] are in use at the Harvard Brigham and Women's Surgical Planning Laboratory as part of a classification procedure for the processing of MR and X-ray CT data. Applying the ideas in this dissertation might provide improvements in VCD image preprocessing.

Finally, the advances in automating VCD suggested earlier based on multiscale statistics can be applied in processing the non-homogeneous images that are common in medical MR imaging. Surface coils are used in MRI when greater intensity resolution is required to illuminate a particular portion of the anatomy. The resulting image has a characteristic non-homogeneous response in intensity. While the human visual system is capable of perceiving a wealth of information in such images, most computer vision algorithms are not designed to handle them. Efforts to correct adjust the gain or intensity amplification across such images have met with some success [Wells 1994]; however, they include preprocessing nonlinear diffusion steps that still require manual parameter selection. Figure 6.2 shows an MR image of a shoulder acquired using a surface coil. A lot of detail is present, but the intensity falloff represents a significant challenge for computer vision. This dissertation provides a foundation for new work in preprocessing MR images acquired from surface coils or exhibiting other non-homogeneous responses. The growth of Functional MRI and its prevalent use of surface coils as well as the advent of new imaging systems employing open-magnets exhibiting non-homogeneous induced magnetic fields will continue to challenge researchers in computer vision.

6.5. Summary

Through this research, I have introduced and explored a new family of statistical operators which demonstrate invariance under rotation, translation, zoom, and linear functions of intensity. These operators are the local expressions of central moments of image intensity, reflecting the distribution of intensities of the neighboring pixel values. Central moments with a directional component have also been constructed using the same principles. Taken as a whole, the combination of local central moments, both isotropic and directional, provides a description of the behavior of local image intensities, capturing changes in both noise and image geometry.

This work is based on the interplay between image intensity and image space. An image is assumed to have piecewise constant behavior with respect some statistical

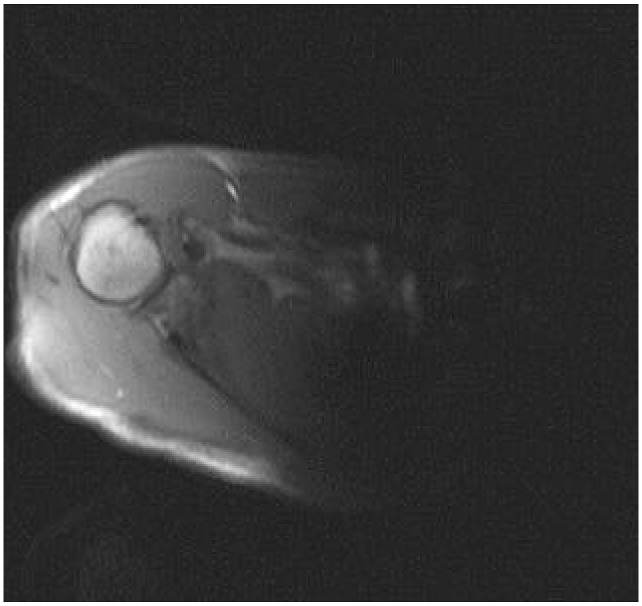

Figure 6.2. MR image of a shoulder acquired using a surface receiving coil. This may represent the ultimate test for this research.

properties of its intensity. For example, these properties include but are not limited to piecewise constant mean, piecewise constant variance, and piecewise mean linearity. These traits are studied by modeling an image as a stochastic process of image space and evaluating its ergodic properties in local areas of the image. Local area image analysis is based on the piecewise ergodic assumption, and trades some regularization or smoothing of local space for insight into image structure.

Although this research has been demonstrated using computer generated images, it has been directed and focused with the goal of applications on real data. Examples of images with varying noise properties and mean intensity are found throughout computer vision. New medical scanning technologies continue to provide computer scientists with challenges in segmentation, classification, object recognition, and image understanding. Market pressures in video communications make issues such as data compression based on object recognition an important research area.

Multiscale image statistics is a new domain, and it is one which has barely been explored. As a basis upon which to decompose images and objects within images, it is somewhat difficult to use since it describes not just the geometry with respect to the image function, but encompasses the uncertainty as well. If data are to be faithfully filtered and represented, some measure of the uncertainty of the locations of the boundaries, the expected geometry, and the noise in that geometry is surely relevant to anyone's analysis.

If we seek answers in object recognition and image understanding, let us understand our assumptions and accurately describe the expected errors in our solutions. This dissertation has been an exploration of new domains in statistics in computer vision to provide fulcrums for our levers, to make approachable the realm of evaluating error in object measurement.